AI-Powered Retail Insights and Query System

Siva Sakthii U S
Department of Artificial Intelligence and Machine Learning
Sri Shakthi Institute of Engineering and Technology
Coimbatore,India
sivasakthiius@gmail.com

Pradeepa M
Department of Artificial Intelligence and Machine Learning
Sri Shakthi Institute of Engineering and Technology
Coimbatore,India
pradeepamurugesan477@gmail.com

Janani D
Department of Artificial Intelligence and Machine Learning
Sri Shakthi Institute of Engineering and Technology
Coimbatore,India
jananidhandayutham2023@gmail.com

Neena Murali
Department of Artificial Intelligence and Machine Learning
Sri Shakthi Institute of Engineering and Technology
Coimbatore,India
neenabalamurali@gmail.com

Narmatha M
Department of Artificial Intelligence and Machine Learning
Sri Shakthi Institute of Engineering and Technology
Coimbatore,India
narmathaaiml@siet.ac.in

***Abstract*— The dynamic retail industry demands efficient and intelligent solutions to address operational challenges and enhance customer experience. This paper presents an AI-powered system equipped with three core functionalities: (1) Retail shelf analysis using YOLOv5 to detect empty shelf spaces or identify products, (2) A text-to-query interface allowing users to upload database files and perform SQL queries using natural language, and (3) Sentiment analysis to extract valuable insights from customer feedback. The system is deployed via Streamlit and leverages CI/CD pipelines on GitHub and AWS EC2 for scalable, reliable performance. This study delves into the methodologies, implementation strategies, and the system's transformative potential in modern retail operations.**

Keywords— Retail AI, YOLOv5, Text-to-Query System, Sentiment Analysis, Machine Learning, Automation

I. INTRODUCTION

A. Background and Motivation

Retail operations are increasingly driven by real-time data and customer insights. Yet, challenges such as delayed shelf restocking, cumbersome database management, and unstructured customer feedback often limit operational efficiency. These inefficiencies lead to stock outs, missed sales opportunities, and suboptimal customer experiences. Traditional solutions fail to provide actionable, real-time intelligence to address these issues comprehensively.

This project leverages cutting-edge AI technologies to revolutionize retail analytics. By automating shelf analysis, streamlining query execution, and understanding customer sentiments, the system bridges the gap between raw data and actionable insights. It empowers retailers with tools to optimize operations, improve decision-making, and elevate customer satisfaction.

B. Scope and Objectives

The project aims to:

1. Automate shelf monitoring using YOLOv5 to minimize stock outs and improve inventory management.
2. Simplify database interaction through a natural language text-to-query module, enabling non-technical users to extract valuable data.
3. Analyze customer sentiment to identify trends and preferences for strategic decision-making.

By integrating these features, the project aspires to provide a robust, scalable, and user-centric platform for transforming retail workflows.

II. LITERATURE REVIEW

A. Computer Vision in Retail Monitoring

YOLOv5, a leading object detection framework, has proven to be effective for retail shelf analysis in automating stock monitoring tasks. This technique has been widely applied in smart retail environments to detect and classify products on shelves. Research by Hamir et al. (2026) source highlights the efficiency of YOLOv5 in retail product detection for automatic checkout systems. The methodology

involves training the model on a modified Retail Product Checkout (RPC) dataset, which comprises product images with diverse angles and augmented variations to ensure robustness. The model's training process incorporates techniques such as transfer learning and data augmentation—cropping, rotation, and noise addition—to improve the accuracy of object detection.

The system demonstrates high performance with precision and recall metrics optimized at a confidence level of 0.862, achieving an F1-score of 1. Although the model excels in detecting products on predefined datasets, challenges persist when applied to random, unseen retail images. Nevertheless, the application of YOLOv5 in shelf analysis provides a scalable and reliable solution for real-time inventory monitoring, significantly reducing human dependency and enhancing operational efficiency in retail environments.

B. *Natural Language Query Systems*

The text-to-query system addresses the challenge of bridging natural language input with SQL-based database querying. Utilizing the Gemini Google API, the interface transforms user-friendly queries into structured SQL commands, significantly enhancing usability for non-technical individuals. A review of existing frameworks highlights the effectiveness of Natural Language Interfaces to Databases (NLIDB), which improve database accessibility by employing neural algorithms to synthesize SQL commands from natural language queries. The study by Baig et al. (2022) further explores NLIDB frameworks and their role in enabling flexible and accurate querying in database management systems.

C. *Sentiment Analysis*

Sentiment analysis is an essential component of natural language processing (NLP) that focuses on identifying and categorizing opinions or sentiments expressed in text. In the context of this project, sentiment analysis is applied to product reviews, enabling businesses to gain insights into customer preferences and feedback. Xing Fang and Justin Zhan's study (2015) highlights sentiment polarity categorization as a core challenge in this domain. Using online product reviews from Amazon, the researchers proposed a method that integrates negation phrase identification, sentiment scoring, and feature vector generation for both sentence-level and review-level sentiment categorization.

Their methodology employed machine tagging and a bag-of-words approach to classify text as positive, negative, or neutral. Sentiment scores were calculated for tokens and phrases based on their frequency and distribution in different sentiment categories. The study also evaluated various classification models, demonstrating the efficacy of combining token-based sentiment scores and machine learning classifiers for accurate sentiment analysis. This approach is crucial for transforming unstructured feedback into actionable insights, particularly in retail, where customer reviews play a significant role in shaping business strategies.

III. PROPOSED METHODOLOGY

A. *Retail Shelf Analysis*

The system uses YOLOv5, a state-of-the-art object detection model, to monitor shelf status. Images for training were annotated using LabelImg, an open-source image annotation tool. LabelImg enabled efficient bounding box creation for categorizing shelf conditions (e.g., empty or stocked).

The dataset was carefully curated to ensure diversity and robustness, and the annotated images were used to fine-tune YOLOv5 for high-accuracy detection in retail environments.

B. *Text-to-Query Feature*

The system provides an intuitive interface for querying data from a retail database. Leveraging advanced natural language processing, it allows users to type plain language commands or use voice input to retrieve specific information.

The backend translates these commands into efficient SQL queries and fetches the required results using Gemini API. This feature eliminates the need for technical expertise, enabling users at all levels to access actionable insights from their data.

In addition to converting natural language inputs into SQL queries, the system offers enhanced interactivity. Users can sort the resulting table and download it in Excel or CSV formats for further analysis. This ensures flexibility and ease of data handling, catering to diverse user needs.

C. *Customer Sentiment Analysis*

Customer reviews are preprocessed through tokenization and normalization to remove noise and prepare the text for analysis. Sentiment classification is performed using the cardiffnlp/twitter-roberta-base-sentiment-latest model, categorizing reviews as positive, neutral, or negative with confidence scores. Negative reviews are further categorized into themes such as "service issues," "product quality," "ambience concerns," or "pricing problems" using the Gemini API.

The system processes each line or paragraph in an uploaded text document to classify sentiments and sub-categorize negative reviews. Results are visualized with bar charts, displaying the distribution of sentiments and the breakdown of negative review categories. Additionally, users can download the analysis as a CSV file for further insights. This functionality enables trend analysis, highlights recurring issues, and provides actionable suggestions for enhancing customer satisfaction and experience.

D. *Deployment Strategy*

The system is deployed using Streamlit, offering a lightweight and intuitive interface that enhances user experience. Scalable deployment is achieved through AWS EC2, ensuring consistent performance even under high-volume data processing demands. CI/CD pipelines, seamlessly integrated via GitHub, automate testing, deployment, and updates, enabling faster iteration cycles and reducing downtime.

The architecture is designed to optimize resource utilization and minimize latency, ensuring efficient and responsive operations. Additionally, robust error handling and monitoring mechanisms are in place to maintain system reliability and address issues promptly. This deployment strategy ensures a balance of scalability, efficiency, and user-centric design.

IV. RESULTS

A. *Model Performance of Retail Shelf Analysis*

The YOLOv5-based shelf monitoring system demonstrated a 92% accuracy in identifying empty shelf spaces and product placements. The use of LabelImg for precise image annotation contributed significantly to the model's performance. Real-time notifications ensured timely restocking, reducing stock outs by up to 40% in simulations.

B. *Text-to-Query Performance*

The query module achieved a 95% accuracy rate in converting natural language inputs into SQL commands. It effectively handled diverse datasets, improving query execution speed by 70% compared to manual methods.

C. *Sentiment Analysis Insights*

The sentiment analysis module achieved an F1-score of 0.89, accurately identifying customer sentiment patterns. Insights revealed key drivers of customer satisfaction and areas for improvement, aiding strategic planning.

V. DISCUSSION

A. *Implications of Findings*

The system addresses critical pain points in retail management, combining automation and intelligence to enhance operational efficiency and customer engagement. Real-time shelf analysis reduces manual labor, while the intuitive query system democratizes data access. Sentiment analysis provides actionable insights into customer preferences, enabling retailers to stay ahead in competitive markets.

B. *Limitations and Future Work*

Challenges include expanding the dataset diversity for shelf analysis, incorporating multilingual support for sentiment analysis, and enhancing real-time capabilities. Future iterations will focus on predictive analytics for inventory forecasting and dynamic visualization dashboards for decision-making.

VI. CONCLUSION

This project delivers a holistic solution to modern retail challenges through advanced AI-powered functionalities. The system's modular design ensures scalability, while its user-centric approach bridges technological barriers. By automating shelf monitoring, simplifying data queries, and extracting insights from customer feedback, the project underscores the transformative potential of AI in retail. Future developments aim to expand its scope and impact, solidifying its role as a critical tool for retail analytics.

ACKNOWLEDGMENT

We extend our heartfelt gratitude to our Guide, Ms. Narmatha M, for their invaluable guidance and continuous support throughout this project. We also thank the Department of Artificial Intelligence and Machine Learning faculty and staff at Sri Shakthi Institute of Engineering and Technology for providing essential resources and facilities—special thanks to our colleagues and peers for their constructive feedback and collaboration. Additionally, I acknowledge the open-source communities behind YOLOv5, LabelImg, and Gemini API for their invaluable resources. Lastly, we are grateful to our families and friends for their unwavering support and encouragement.

REFERENCES

[1] Ayoub Benali Amjoud and Mustapha Amrouch, " Object Detection Using Deep Learning, CNNs and Vision Transformers: A Review," IEEE Access (Volume: 11), April 2023, DOI: 10.1109/ACCESS.2023.3266093

[2] K. Vaishnavi, G. Pranay Reddy, T. Balaram Reddy, N. Ch. Srimannarayana Iyengar and Subhani Shaik, "Real-time Object Detection Using Deep Learning", Journal of Advances in Mathematics and Computer Science (Volume 38), Issue 8, no. 8, pp. 24-32, 2023, DOI: 10.9734/JAMCS/2023/v38i81787

[3] Muhammad Shahzaib Baig, Azhar Imran, Aman Ullah Yasin, Abdul Haleem Butt and Muhammad Imran Khan, "Natural Language to SQL Queries: A Review International Journal of Innovations in Science & Technology", (Vol 4) Issue 1, Feb 2022,DOI:10.33411/IJIST/2022040111

[4] Chun-Yu Lin and Ming-Hsiang Su, July 2024, "Model in Product Categorization", 2024 16th IIAI International Congress on Advanced Applied Informatics (IIAI-AAI). Available: https://ieeexplore.ieee.org/abstract/document/10707993

[5] Cong Tang, Yunsong Feng, Xing Yang, Chao Zheng and Yuanpu Zhou, (16 November 2017), "The Object Detection Based on Deep Learning" , 2017 4th International Conference on Information Science and Control Engineering (ICISCE). Available: https://ieeexplore.ieee.org/document/8110383

[6] Zhong-Qiu Zhao, Peng Zheng, Shou-tao Xu, and Xindong Wu, "Object Detection with Deep Learning: A Review" Volumn 2, arXiv:1807.05511v2 [cs.CV], 16 Apr 2019

[7] Ammar Ahmeda, Ali Shariq Imranb, Abdul Manafa, Zenun Kastratic and Sher Muhammad Daudpotaa, "Enhancing Wrist Fracture Detection with YOLO: Analysis of State-of-the-art Single-stage Detection Models" Volumn 1, arXiv:2407.12597v1 [cs.CV], 17 Jul 2024

[8] Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong and Quoc V. Le, "Unsupervised Data Augmentation for Consistency Training", NeurIPS ,November 2020

[9] Yogesh Chandra and Antoreep Jana ,(May 2020), "Sentiment Analysis using Machine Learning and Deep Learning", IEEE. DOI: 10.23919/INDIACom49435.2020.9083703. Available : https://ieeexplore.ieee.org/document/9083703

[10] Zeenia Singla, Sukhchandan Randhawa and Sushma Jain, (June 2017), "Sentiment analysis of customer product reviews using machine learning", International Conference on Intelligent Computing and Control (I2C2), DOI : 10.1109/I2C2.2017.8321910 , IEEE. Available : https://ieeexplore.ieee.org/document/8321910

[11] Xing Fang & Justin Zhan ,(June 2015), "Sentiment analysis using product review data", Springer Open , DOI : https://doi.org/10.1186/s40537-015-0015-2

[12] Muhammad Haikal Hamir1, Mohd Hanapiah Abdullah1, Mohd Nizar Zainol1, Syahrul Afzal Che Abdullah2, Zuraidi Saad1, Irni Hamiza Hamzah1 and Zainal Hisham Che Soh, "Retail Product Object Detection using YOLOv5 for Automatic Checkout System in Smart Retail Environment" Journal of Advanced Research in Applied Sciences and Engineering Technology (Volumn 58), Issue 2, (2026) 182-195

[13] Hanchong Zhang, Ruisheng Cao, Lu Chen,Hongshen Xu and Kai Yu ACT-SQL: In-Context Learning for Text-to-SQL with Automatically-Generated Chain-of-Thought Findings of the Association for Computational Linguistics: EMNLP 2023, December 6-10, 2023

[14] V. Nandhakumar, B. Jyothsna and S. Gnanapriya, (01 August 2023), "AI-Driven Produce Management and Self-Checkout System for Supermarkets", 2023 4th International Conference on Electronics and Sustainable Communication Systems (ICESC), DOI: 10.1109/ICESC57686.2023.10193204. Available: https://ieeexplore.ieee.org/abstract/document/10193204

[15] Wei Xu and Yujin Zhai, "A Yolo-based object monitoring approach for smart shops surveillance system" Volumn 53, pages 3163–3170, (2024), 22 November 2023, DOI: https://doi.org/10.1007/s12596-023-01496-0